

UNITED STATES PATENT APPLICATION FOR:

METHOD AND APPARATUS FOR DETECTING OBSTACLES

INVENTORS:

JOHN RICHARD FIELDS
JOHN BENJAMIN SOUTHALL

ATTORNEY DOCKET NUMBER: SAR 14882

CERTIFICATION OF MAILING UNDER 37 C.F.R. 1.10

I hereby certify that this New Application and the documents referred to as enclosed therein are being deposited with the United States Postal Service on 7/16/03, in an envelope marked as "Express Mail United States Postal Service", Mailing Label No. E470250024145, addressed to: Mail Stop Patent Application, Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.

Kimberly Boyd
Signature

Kimberly Boyd
Name

7/16/03
Date of signature

MOSER, PATTERSON & SHERIDAN LLP
595 Shrewsbury Ave.
Shrewsbury, New Jersey 07702
(732) 530-9404

METHOD AND APPARATUS FOR DETECTING OBSTACLES

GOVERNMENT RIGHTS IN THIS INVENTION

[0001] This invention was made with U.S. government support under contract number MDA972-01-9-0016. The U.S. government has certain rights in this invention.

BACKGROUND OF THE INVENTION

Field of the Invention

[0002] The invention relates to vision systems and, more particularly, the present invention relates to a method and apparatus for detecting obstacles using a vehicular-based vision system.

Description of the Background Art

[0003] Vehicular vision systems generally comprise a camera (or other sensor) mounted to a vehicle. An image processor processes the imagery from the camera to identify obstacles that may impede the movement of the vehicle. To identify obstacles, a plane is used to model the roadway in front of the vehicle and the image processor renders obstacles as point clouds that extend out of the plane of the roadway. By using such a planar model, the processing of imagery from "on-road" applications of vehicular vision systems is rather simple. The image-processing system must recognize when the point cloud is extending from the roadway plane and deem the point cloud simply to be an obstacle to be avoided.

[0004] In "off-road" applications, where the ground upon which the vehicle is to traverse is non-planar, the terrain cannot be modeled as a simple plane. Some applications have attempted to model the off-road terrain as a plurality of interconnecting planes. However, such models are generally inaccurate and cause the vehicle to identify obstacles that could, in reality, be traversed by the vehicle. As such, unnecessary evasive action is taken by the vehicle.

[0005] Therefore, there is a need for a method and apparatus of performing improved obstacle detection that is especially useful in "off-road" applications.

SUMMARY OF THE INVENTION

[0006] The invention provides a method and apparatus for detecting obstacles in non-uniform environments, e.g., an off-road terrain application. The apparatus uses a stereo camera and specific image-processing techniques to enable the vehicle's vision system to identify drivable terrain in front of the vehicle. The method uses the concept of a non-drivable residual (NDR), where the NDR is zero for all terrain that can be easily traversed by the vehicle and is greater than zero for terrain that may not be traversable by the vehicle. The method utilizes a depth map having a point cloud that represents the depth to objects within the field of view of the stereo cameras. The depth map is organized into small tiles; each tile is represented by the average of the point cloud data contained within. The method scans columns of pixels in the image to find sequences of "good" points that are connected by line segments having an acceptable slope. Points that lie outside of the acceptable slope range will have an NDR that is greater than zero. From this information regarding obstacles and the terrain before the vehicle, the vehicle control system can accurately make decisions as to the trajectory of the vehicle.

BRIEF DESCRIPTION OF THE DRAWINGS

[0007] So that the manner in which the above recited features of the present invention are attained and can be understood in detail, a more particular description of the invention, briefly summarized above, may be had by reference to the embodiments thereof which are illustrated in the appended drawings.

[0008] It is to be noted, however, that the appended drawings illustrate only typical embodiments of this invention and are therefore not to be considered limiting of its scope, for the invention may admit to other equally effective embodiments.

[0009] Figure 1A and 1B depict a vehicle on off-road terrain;

[0010] Figure 2 depicts a block diagram of a vision system in accordance with the present invention;

[0011] Figure 3 depicts a functional block diagram of various components of the vision system in accordance with the present invention;

- [0012] Figure 4 depicts a flow diagram of a process of operation of the present invention;
- [0013] Figure 5 depicts a terrain model with decision points suggested by the present invention; and
- [0014] Figure 6 depicts one column of depth map data as processed by the present invention.

DETAILED DESCRIPTION

[0015] Figure 1A depicts a side view of a vehicle 100 having a movement system 101 traversing off-road terrain and Figure 1B depicts a top view of the terrain in Figure 1A. The vehicle 100 contains a stereo imaging system 102 having at least a pair of sensors or cameras mounted to the front of the vehicle. In one illustrative embodiment, the vision system 102 is capable of processing video at a rate of ten frames per second or faster in real time and produces an obstacle map that has a resolution that is fine enough to identify a pathway that is a little wider than the vehicle itself. The vehicle may be an unmanned ground vehicle (UGV) that uses the obstacle detection method of the present invention to enable the vehicle's control system to direct the vehicle around detected obstacles. Alternatively, the invention could also be used as an obstacle avoidance warning system for a manned vehicle or for a system that detects the slope of terrain to enable a driver to understand whether the slope is traversable by the vehicle without causing damage to the vehicle.

[0016] The method does not recognize specific objects but labels areas that are difficult or impossible to traverse. Also, the method does not determine if an area on the other side of an obstacle can be reached, that process is left to the route planner that is responsible for that task.

[0017] An advantage of the non-drivable residual (NDR) method of the present invention is that it enables the evaluation of the change in vertical height from one place to another relative to the range of heights that would occur for drivable slopes. As such, the method uses mobility constraints for a particular vehicle and compares the constraints to the slope of the terrain proximate the vehicle. As illustrated in Figure 1A, the process starts from a "good" point 106 that lies on the surface of the terrain in

front of the vehicle. A "good" point is a surface point on terrain that can be traversed by the vehicle, e.g., fulfills a mobility constraint. The stereo images captured by the system 102 are converted into a depth map that shows the depth of points within the field of view of the vehicle's cameras. By processing the images, a depth map can be used to identify the terrain profile 108. Each point in the profile 108 is compared to the good point 106 and the non-drivable residual indicates a departure between the height of the next point along the profile and the interval of heights that could be reached if the same distance were traversed on a drivable slope. As long as the height of the next point lies within the drivable range as indicated by the boundaries 110 and 112, the residual is zero, and the "good" point is updated accordingly. The residual becomes non-zero when the height exceeds the drivable range outside of the boundaries 110 and 112, i.e., above the point 114 on the terrain profile 108. The "good" point becomes fixed, and subsequent points are evaluated relative to this reference. The residual itself is measured from the appropriate limiting slope line 110. If the non-drivable residual exceeds a threshold, then an impassible obstacle has been detected, e.g., the mobility constraint is exceeded for the particular vehicle. In the example shown in Figures 1A and 1B, the vision system will deem the obstacle 104 non-traversable by the vehicle. Further examples will be discussed below as the hardware and software of the present invention are described.

[0018] Figure 2 depicts a block diagram of one embodiment of the hardware that can be used to form the vision system 102. The vision system 102 comprises a pair of charge coupled device (CCD) cameras 200 and 202 that form a stereo imaging system. The vision system 102 further comprises an image processor 204 that is coupled to the cameras 200 and 202. The image processor 204 comprises an image preprocessor 206, a central processing unit (CPU) 208, and support circuits 210 and memory 212. The image preprocessor 206 comprises circuitry that is capable of calibrating, capturing and digitizing the stereo images from the cameras 200 and 202. Such an image preprocessor is the Acadia integrated circuit available from Pyramid Vision Technologies of Princeton, New Jersey. The central processing unit 208 is a general-purpose computer or microprocessor. Support circuits 210 are well known and are used to support the operation of the CPU 208. These circuits include such well-known circuitry as cache, power supplies, clock circuits, input/output circuitry and the like. Memory 212 is coupled to the CPU 208 for storing a database, an operating

system and image processing software 214. The image processing software 214, when executed by the CPU 208, forms part of the present invention.

[0019] Figure 3 depicts a functional block diagram of the various modules that make up the vision system 102. The cameras 200 and 202 are coupled to the stereo image preprocessor 300 that produces stereo imagery. The stereo imagery is processed by the depth map generator 302 to produce a depth map of the scene in front of the vehicle. The depth map comprises a two-dimensional array of pixels, where a value of a pixel represents the depth to a point in the scene. The depth map is processed by the depth map processor 304 to perform piecewise smoothing of the depth map and identify obstacles within the path of the vehicle. The obstacle's detection information is coupled to the vehicle controller 306 such that the vehicle controller can take action to avoid the obstacle, warn a driver, plan and execute an optimal route, and the like.

[0020] Figure 4 depicts a method 400 of operation of the vision system illustrated in Figures 1-3. The method 400 begins, at step 402, by producing a depth map of the scene in front of the vehicle. Generally, this is accomplished by the Acadia circuitry. At step 404, the depth map is then piecewise smoothed. The smoothing is performed by dividing the depth map into small portions; e.g., 5 pixel by 5 pixel blocks. The planar tile is fit to the pixels in each of the blocks. The center of each tile is used as a "point" in processing the depth map. Then, at step 406, an initial last point and an initial last good point are established. These initial values can be default values or values determined by the particular scene. At step 408, a current point is selected within the smoothed depth map. The method 400 generally processes the smoothed depth map by selecting a point within a selected column of points, processing all the points in a column and then processing the next adjacent column of points and so on. Alternatively, a row of points across all columns can be processed simultaneously, and then each higher row of points is processed until all the points in the smoothed depth map are processed. To ensure accuracy, the points identified as good can be compared in rows of points to ensure consistency or to compensate for data drop-outs.

[0021] At step 410, the method 400 determines whether the current point is within the drivable slope of the last point. If the answer is negative, the method 400 proceeds to step 412 and determines if the current point is within the drivable slope of the last good point. If the current point is within the drivable slope, or if the current point was

determined in step 410 to be within the drivable slope of the last point, at step 414, the NDR for the current point is set to 0, and the last good point is updated to the current point. Then, at step 416, the zero NDR for the current point is stored. At step 418, the method 400 determines whether there is more data to be processed. If there is more data, the last point is updated to the current point at step 420 and a loop is made back to step 408 for the selection of a new current point.

[0022] However, if during step 412 a determination is made that the current point is not within the drivable slope of the last good point, a non-zero NDR with respect to the last good point is calculated for the current point at step 424. At step 416, the non-zero NDR is then stored for the current point. At step 418, the method queries whether more data is to be processed. If the query is affirmatively answered, the method 400 proceeds to step 420 and sets the current point to the last point and proceeds to step 408 to process the next point.

[0023] When a determination is made in step 418 that there is no more data to be processed, the method 400 proceeds to step 426 wherein the points and NDRs are projected onto a map. At step 428, the map is used to plan a route that will avoid any detected obstacle. The plan may then be executed. For example, the map contains a two-dimensional matrix of values where zero value and low values represent passable terrain (i.e., terrain that does not exceed the mobility constraint of the vehicle) and high values represent impassable terrain (i.e., terrain that exceeds the mobility constraint of the vehicle). The specific thresholds assigned to produce "low" and "high" indications are defined by the particular vehicle that is traversing the terrain. Consequently, the map identifies regions in which the mobility constraints of the particular vehicle are exceeded and not exceeded.

[0024] Figure 5 depicts a schematic view of various points along a slope as processed by the method of Figure 4. The first point 502 is assumed to be "good". The second point 504 is within the drivable interval of the boundaries extending from the first point, as such, point 504 is deemed good. The third point 506 is outside the drivable interval of the second point, and its residual is calculated as discussed below. The fourth point 508 is also outside the drivable interval of the third point 506, and its residual is computed with respect to the drivable interval of the second point 504. This "frozen" last good point 504 becomes a fixed local reference for evaluating the severity

of a potential obstacle. The obstacle ends with point 510 since the elevation is within the drivable range of the previous point 508. In this example the vehicle would easily traverse the terrain through points 502 and 504; however, the NDR of point 506 would be evaluated to see if it is above the threshold for the vehicle to traverse the terrain at that angle. The same is true for the terrain at point 508. If the NDR is severe enough, then the method 400 will deem the terrain at points 506 and 508 to be non-drivable. However, if the NDR is not substantial, then the terrain feature (such as a small rock) is considered to be passable, even though the slope is outside of the boundaries extrapolated from point 504.

[0025] Figure 6 illustrates how the method of the present invention operates on data representing a large rock on a small incline. The camera viewing the scene is located to the left of the figure. The lines 602 indicate drivable slopes. The diamonds are points that are mapped into pixels in one column of the smoothed depth map. The first 6 points are "good" points. The next point is outside of the limits of the boundaries but is not far enough from the drivable slope to be classified as an obstacle. Points 8 through 11 exceed a threshold and are classified as obstacles. The first point visible above the rock is again a good point as are points 13 and 14. The table to the right of the figure lists the non-drivable residual for each point. The threshold in this case is set at 0.1.

[0026] The following calculation is applied to pixels (points) in one column of the image at a time. If there is a stereo dropout (unavailable data), the computation continues with the next available pixel. The only state variables are the last point and the last good point. As mentioned above, the points may be processed simultaneously in rows and further comparative processing can be performed to ensure accuracy of the computations.

[0027] Let (X, Y, Z) be the world coordinates of the point imaged at pixel (x, y) in the image. Assume that the world coordinates have been suitably transformed so that the Y axis is vertical. In practice, this transformation is achieved with input from an inertial navigation system (INS) which relates the camera pose to the world system. (In the usual system, X points right, Y points down, and Z points forward.) Let $(X, Y, Z)_L$ be the coordinates of the last point, and $(X, Y, Z)_G$ be the coordinates of the last "good" point. The initial values of these points are:

$$(X, Y, Z)_L = (X, Y, Z)_G = (0, -h, 0)$$

where h is the camera height.

[0028] To compute the non-drivable residual (NDR or R_{nd}) for point (X, Y, Z) , first compute the displacement from the last point:

$$(\Delta X, \Delta Y, \Delta Z)_L = (X, Y, Z) - (X, Y, Z)_L$$

The distance traveled (projected onto the XZ plane) is:

$$d_L = \sqrt{\Delta X_L^2 + \Delta Z_L^2}$$

Let s_{di} be the maximum slope of a drivable incline (uphill or downhill). The limiting values for a drivable ΔY are:

$$\Delta Y_{uphill} = -s_{di} d_L \text{ and } \Delta Y_{downhill} = s_{di} d_L$$

If $\Delta Y_{uphill} \leq \Delta Y \leq \Delta Y_{downhill}$, then the method has found a nominally flat, level place. Set $R_{nd} = 0$ and update the last point and the last good points:

$$(X, Y, Z)_L \leftarrow (X, Y, Z) \text{ and } (X, Y, Z)_G \leftarrow (X, Y, Z).$$

Otherwise, the change in elevation indicates a possible obstacle. To measure the severity of the height change, first the method computes the distance from the last good point:

$$d_G = \sqrt{\Delta X_G^2 + \Delta Z_G^2} \text{ where } (\Delta X, \Delta Y, \Delta Z)_G = (X, Y, Z) - (X, Y, Z)_G$$

The ΔY limits for computing the residual are:

$$\Delta Y_{uphill} = -s_{di} d_G \text{ and } \Delta Y_{downhill} = s_{di} d_G$$

The residual is given by:

$$R_{nd} = \begin{cases} \Delta Y_G - \Delta Y_{downhill}, & \Delta Y_{downhill} < \Delta Y_G \\ 0 & \Delta Y_{uphill} \leq \Delta Y_G \leq \Delta Y_{downhill} \\ \Delta Y_G - \Delta Y_{uphill} & \Delta Y_G < \Delta Y_{uphill} \end{cases}$$

The residual is compared to a pre-defined threshold. If the residual is greater than the threshold, then the potential obstacle is deemed an actual obstacle to be avoided, i.e., the terrain is not traversable. Lastly, the method always updates the last point:

$(X, Y, Z)_L \leftarrow (X, Y, Z)$ and, if $R_{nd} = 0$, then the method also updates the last good point:

$(X, Y, Z)_G \leftarrow (X, Y, Z)$.

[0029] Spurious values in the obstacle map can be suppressed by applying the method to average values of (X, Y, Z) . Most of the experiments and tests have been done with averages computed for non-overlapping blocks of 5x5 pixels. Good results have also been obtained for overlapping, variable, sized patches ranging from 40 pixels square in the foreground to a minimum of 8 pixels square at row 68 out of 320. The main issue with the larger, overlapping averages is the increase in computation time. To obtain average values of (X, Y, Z) , the quantity $(1/Z)$ is approximated by a linear function of the pixel coordinates (x, y) in the patch. The value of $(1/Z)$ obtained from the fit is used to compute Z at the center of the patch. X and Y are then computed from Z , the pixel coordinates, and the camera center and focal length.

[0030] The average is computed in camera coordinates, and then transformed to world coordinates. The transformation matrix includes the camera-to-vehicle rotation obtained from camera calibration, and the vehicle-to-world transformation obtained from the vehicle pose sensors.

[0031] While foregoing is directed to various embodiments of the present invention, other and further embodiments of the invention may be devised without departing from the basic scope thereof, and the scope thereof is determined by the claims that follow.